First Hit

Generate Collection Print

L3: Entry 1 of 2 File: PGPB Dec 26, 2002

DOCUMENT-IDENTIFIER: US 20020198863 A1
TITLE: Stratified sampling of data in a database system

Summary of Invention Paragraph:
[0005] The entire population of data contained in a data set may not be homogenous. For example, in maintaining records of shoppers at a retail outlet, it may be determined that 80% of the shoppers are male while 20% of the shoppers are female. If this is the case, it is sometimes desirable to obtain stratified random samples, as compared to simple random samples. Stratified random sampling involves dividing a given population into homogenous subgroups and then taking a simple random sample in each subgroup. Thus, in the above example, the population is divided into two subgroups, one female and one male.

Brief Description of Drawings Paragraph:
[0010] FIG. 2 is a flow diagram of a process performed by the database system, in accordance with an embodiment, to perform stratified random sampling.

Brief Description of Drawings Paragraph:
[0011] FIGS. 3 and 4 illustrate processes of performing random sampling in each stratum.

Detail Description Paragraph:
[0017] The arrangement of the database system 14 shown in FIG. 1 is an example of a parallel database arrangement, in which the multiple AMPs 24 are capable of concurrently accessing and manipulating data in respective storage modules 26. Each relational table stored in the database system is partitioned across the multiple AMPs and respective storage modules. In other words, a given table is divided into multiple partitions and stored in respective storage modules. In other embodiments, instead of a parallel database system, a single-node or uni-processor database system can be used.

Detail Description Paragraph:
[0018] An efficient technique is provided to perform stratified random sampling in the database system 14. In stratified random sampling, a data set containing a given population of data is divided into multiple subgroups (or strata). Within each subgroup (or stratum), random sampling is performed. Thus, for example, a population may be divided into male and female subgroups. Stratification can also be performed on the basis of age, professions, or other criteria.

Detail Description Paragraph:
[0019] In accordance with some embodiments of the invention, to enable efficient stratified random sampling, an SQL query (or a query according to another standard database query language) is extended to add a predefined clause (referred to as the SAMPLE STRATIFIED clause in one example embodiment) to indicate stratified sampling is to be performed. Note that the predefined clause to indicate performance of stratified sampling can have other names.

Detail Description Paragraph:
[0023] Note that the SAMPLE STRATIFIED clause in the SELECT query enables multiple stratification conditions to be specified in one query. This makes stratified

random sampling more efficient as multiple separate queries need not be submitted to perform the stratified random sampling.

Detail Description Paragraph:
[0026] FIG. 2 illustrates a process according to one embodiment of the invention for performing stratified random sampling. Upon receiving a query, the parsing engine 20 parses and analyzes the query (at 102). The parsing engine 20 looks for the SAMPLE STRATIFIED clause to determine (at 104) if the query is a stratified random sampling query. If not, then other processing is performed (at 106). However, if the query is a stratified random sampling query, then the parsing engine 20 generates (at 108) an evaluation plan that includes the creation of input spool files 150 (FIG. 1) for stratified sampling. The parsing engine 20 allocates (at 112) the input spool files to store random samples, with one spool file allocated per stratum. Thus, if the query specifies 3 strata, then 3 corresponding spool files are created.

Detail Description Paragraph:
[0030] The following provides some examples of simple and complex queries for purposes of illustration. A WHERE clause of a simple SQL query involves a selection criterion that involves columns from one table only. Thus, the example query, SELECT Name FROM stratified random sampling query. However, if the WHERE clause of the query involves joins, then the query is a complex query. An example of a complex query with a join clause is:

Detail Description Paragraph:
[0034] In response to the step(s) received at 119, each AMP evaluates (at 120) the query condition to obtain qualifying records. Based on the stratification conditions (contained either in an enhanced step or a separate stratification step), the AMP writes qualifying records to corresponding spool files. Next, beginning with the first stratum (at 124), each AMP performs the random sampling of records of each input spool file (at 126). The random sampling according to one embodiment is described in connection with FIGS. 3 and 4 below. Each AMP then determines (at 128) if there are more strata left. If so, the next stratum is processed (at 130). This is repeated until all strata have been processed and the random samples have been collected in each stratum.

Detail Description Paragraph:
[0037] As shown in FIG. 3, a general process of obtaining random samples in plural strata is illustrated. The AMP receives N input spool files 202, 208 that correspond to plural strata. Random sampling is then performed (at 204, 210) of records in each spool file. Each AMP uses a pseudo-random number generator 152 (FIG. 1) to perform the random sampling. The sampling algorithm employed is random so that any sample data in a table is equally likely to be selected. The rows obtained as a result of the random sampling of each spool file are outputted as sample rows in a corresponding output file 206, 212.

Detail Description Paragraph:
[0038] If the database system 14 is a single uni-processor system with a single AMP, then the random sampling for each stratum is relatively straightforward, as the spool file for the stratum is non-partitioned. However, in a parallel database system environment where each spool file is partitioned across plural AMPs and stored in a plurality of partitions in respective storage modules, the random sampling is performed according to a process in FIG. 4 in one embodiment.

Detail Description Paragraph:
[0039] To perform random sampling in a parallel database system environment, the number of sample rows to be obtained at each AMP are first pre-allocated. The parallel random sampling algorithm preserves most of the randomness properties as the pre-allocation does not examine the data itself. Within each stratum, the input to the parallel random sampling process is the input spool file, which is

partitioned across the multiple AMPs of the database system 14.

Detail Description Paragraph:
[0040] FIG. 4 shows a parallel random sampling process performed in accordance with one embodiment. The process is described with respect to one input spool file associated with one stratum. The same process is repeated for other input spool files associated with other strata. Each AMP determines the number of rows in the input spool file partition stored by the AMP by scanning (at 300) an index or the partition to obtain a count of rows stored on the corresponding storage module that is managed by the AMP. In some embodiments, each partition of an input spool file is stored as a B+ Tree indexed file (or some other type of index) on a respective storage module 26. The index contains information of how many rows are in each of the partitions. In one embodiment, the index is scanned to collect row counts for the partition. In other embodiments where no such index is available or where no such information is available in the index, each partition is scanned to obtain the row count. This row count is used in producing random samples of the input spool files generated by stratification.

Detail Description Paragraph:
[0045] The AMP then performs a random "toss" (at 316) to determine if the next row belongs to a sample (it is accepted or rejected). For example, assume that there are N rows stored in an AMP, and that it is desired to select n rows at random from the set of N rows, where 0<n.ltoreq.N. Initially, variables t and m are set to zero (t.rarw.0, m.rarw.0), with m representing the number of rows selected (accepted) so far, and t representing the total number of input rows the AMP has processed. Then, a random number U is generated that is uniformly distributed between 0 and 1.

Detail Description Paragraph:
[0048] The acts are repeated until the requested number of samples have been obtained. If m<n, then the AMP continues the sampling (along the "No" prong of decision block 312). However, if m<n is not true, then the sample is complete and the process terminates. Note that other algorithms can be used for parallel random sampling processes according to other embodiments.

Detail Description Paragraph:
[0049] One potential application of the stratified random sampling technique discussed here is the use of segmentation as a data mining technique. Segmentation includes subdividing a population according to known discriminators for marketing analysis. By using the stratified sampling technique discussed here, segmentation efficiency is enhanced.

CLAIMS:

13. The article of claim 12, wherein the instructions when executed cause the database system to perform random sampling of data in each spool file to obtain samples for a corresponding stratum.

14. The article of claim 13, wherein each spool file is partitioned across the plural access modules, wherein the instructions when executed cause the database system to perform the random sampling of each spool file by performing random sampling in each access module.

15. A database system comprising: a storage to store a base table; and a controller adapted to receive a request containing plural stratification conditions to divide data in the base table into corresponding plural strata, the controller adapted to perform random sampling, in response to the request, of data in each stratum.

18. The database system of claim 17, wherein the controller is adapted to generate plural spool files to store data in the plural strata; and wherein the controller is adapted to perform random sampling of data in each spool file.

21. A database system comprising: a plurality of storage modules; a plurality of access modules to manage respective storage modules; and a parsing engine to receive a stratified sampling query specifying plural stratification conditions, the parsing engine to generate one or more commands to indicate performance of the stratified sampling, the parsing engine to send the one or more commands to the access modules, in response to the one or more commands, each access module to generate plural input spool files corresponding to plural strata, the input spool files to store qualifying rows from a source table, the access module to selectively write a given row into one of the input spool files based on which stratification condition the given row satisfies, each access module to further perform random sampling of the rows in each input spool file.

First Hit Fwd Refs

End of Result Set

Generate Collection Print

L3: Entry 2 of 2 File: USPT May 1, 2001

DOCUMENT-IDENTIFIER: US 6226629 B1
TITLE: Method and apparatus determining and using hash functions and hash values

Detailed Description Text (84):
Unlike UEC, sampling for partitioning keys can be quite efficient. For example sample size calculations can be derived from the application of Chernoff-Hoeffding bounds on tail probabilities of the binomial distribution (see Seshadri and Naughton, following Plaxton, et al. (Seshadri et al, "Sampling Issues in parallel Database Systems," Proc. 3.sup.rd International Conference on Extending Database Technology--EDBT, (Springer-Verlag, 1992), pp. 328-343 and Blelloch et al., "A Comparison of Sorting Algorithms for the Connection Machine CM-2," Proc. Of the Symposium of Parallel Algorithms and Architectures, July 1995, pp. 3-16, each of which is herein incorporated by reference).

Detailed Description Text (91):
(There is an added requirement when multiple estimates are derived from a single random sample, but the computation, although straightforward, will not be given here.) In addition, a sequential scan (`skip sequential`) algorithm, as described in Vitter, "Optimum Algorithms for Two Random Sampling Problems," Proc. 24.sup.th Annual Symposium on the Foundations of Computer Science (IEEE) November 1983, pp. 56-64, which in herein incorporated by reference, can be used for producing more efficient database random samples.

Detailed Description Text (136):
c.sub.1, c.sub.2, d.sub.1, d.sub.2 .di-elect cons.X--For example, c.sub.1, c.sub.2, d.sub.1, d.sub.2 are 32-bit random numbers

Detailed Description Text (276):
10. Seshadri, S. and Naughton, J. F. "Sampling Issues in Parallel Database Systems." Proc. 3rd International Conference on Extending Database Technology--EDBT 92, (Springer-Verlag, 1992), pp. 328-343.

Detailed Description Text (278):
12. Vitter, J. S. "Optimum Algorithms for Two Random Sampling Problems." Proc. 24th Annual Symposium on the Foundations of Computer Science, (IEEE: November 1983), pp. 56-64.

Other Reference Publication (10):
S. Seshadri and Jeffrey F. Naughton, "Sampling Issues in Parallel Database Systems," 3rd International Conference on Extending Database Technology, Springer-Verlag, Mar. 23-27, 1992, pp. 328-343.

Other Reference Publication (11):
Jeffrey Scott Vitter, "Optimum Algorithms for Two Random Sampling Problems," 24th Annual Symposium on Foundations of Computer Science, Nov. 7-9, 1983, pp. 65-75.

Print Selection

Clear | Cancel | Print | Print First Page

Section: Page(s): Print Copy: 1

Select?	Document ID	Section(s)	Page(s)	# Pages to print	Database
☑	20020198863	all	all	11	PGPB,USPT,USOC,EPAB,JPAB,DWPI
☑	6226629	all	all	27	PGPB,USPT,USOC,EPAB,JPAB,DWPI

Building	Room	Printer
cpk2	4c32	gbfrptr

L2: Entry 3 of 6 File: USPT May 13, 2003

DOCUMENT-IDENTIFIER: US 6564221 B1
** See image for Certificate of Correction **
TITLE: Random sampling of rows in a parallel processing database system

Abstract Text (1):
A method, apparatus, and article of manufacture for random sampling of rows stored in a table, wherein the table has a plurality of partitions. A row count is determined for each of the partitions of the table and a total number of rows in the table is determined from the row count for each of the partitions of the table. A proportional allocation of a sample size is computed for each of the partitions based on the row count and the total number of rows. A sample set of rows of the sample size is retrieved from the table, wherein each of the partitions of the table contributes its proportional allocation of rows to the sample set of rows. Preferably, the computer system is a parallel processing database system, wherein each of its processing units manages a partition of the table, and some of the above steps can be performed in parallel by the processing units.

Brief Summary Text (3):
This invention relates in general to database management systems performed by computers, and in particular, to random sampling of rows in a parallel processing database system.

Brief Summary Text (9):
Thus, there is a need in the art for improved random sampling of rows stored on a database system.

Brief Summary Text (11):
The present invention discloses a method, apparatus, and article of manufacture for random sampling of rows stored in a table, wherein the table has a plurality of partitions. A row count is determined for each of the partitions of the table and a total number of rows in the table is determined from the row count for each of the partitions of the table. A proportional allocation of a sample size is computed for each of the partitions based on the row count and the total number of rows. A sample set of rows of the sample size is retrieved from the table, wherein each of the partitions of the table contributes its proportional allocation of rows to the sample set of rows. Preferably, the computer system is a parallel processing database system, wherein each of its processing units manages a partition of the table, and some of the above steps can be performed in parallel by the processing units.

Detailed Description Text (9):
In the preferred embodiment of the present invention, the RDBMS software comprises the Teradata.RTM. product offered by NCR Corporation, and includes one or more Parallel Database Extensions (PDEs) 112, Parsing Engines (PEs) 114, and Access Module Processors (AMPs) 116. These components of the RDBMS software perform the functions necessary to implement the RDBMS and SQL standards, i.e., definition, compilation, interpretation, optimization, database access control, database retrieval, and database update.

Detailed Description Text (20):

The preferred embodiment of the present invention includes both a parallel sampling method and an identification of mutually exclusive rows belonging to the samples. Because the sampling method is random, any sample of the data is equally likely to be selected. In this method, however, a stratified random sampling method is used by pre-determing the number of sample rows to be obtained at each PU 102. This method still preserves most of the randomness properties as the pre-allocation does not examine the data itself. Moreover, because the rows are mutually exclusive, a row belonging to one sample can not belong to any other sample, i.e., the sampling is done without replacement.

Detailed Description Text (49):
In summary, the present invention discloses a method, apparatus, and article of manufacture for random sampling of rows stored in a table, wherein the table has a plurality of partitions. A row count is determined for each of the partitions of the table and a total number of rows in the table is determined from the row count for each of the partitions of the table. A proportional allocation of a sample size is computed for each of the partitions based on the row count and the total number of rows. A sample set of rows of the sample size is retrieved from the table, wherein each of the partitions of the table contributes its proportional allocation of rows to the sample set of rows. Preferably, the computer system is a parallel processing database system, wherein each of its processing units manages a partition of the table, and some of the above steps can be performed in parallel by the processing units.

CLAIMS:

1. A method for random sampling of rows stored in a table in a computer system, wherein the table has a plurality of partitions, the method comprising: (a) determining a row count for each of the partitions of the table; (b) determining a total number of rows in the table from the row count for each of the partitions of the table; (c) computing a proportional allocation of a sample size for each of the partitions based on the row count and the total number of rows; and (d) retrieving a sample set of rows of the sample size from the table, wherein each of the partitions of the table contributes its proportional allocation of rows to the sample set of rows.

4. The method of claim 1, wherein the retrieving (d) step further comprises creating the sample set of rows using a stratified random sampling method.

23. An apparatus for random sampling of rows stored in a table, comprising: (a) a computer system having one or more data storage devices coupled thereto, wherein the data storage devices store at least one table, and the table has a plurality of partitions; (b) logic, performed by the computer system, for (1) determining a row count for each of the partitions of the table; (2) determining a total number of rows in the table from the row count for each of the partitions of the table; (3) computing a proportional allocation of a sample size for each of the partitions based on the row count and the total number of rows; and (4) retrieving a sample set of rows of the sample size from the table, wherein each of the partitions of the table contributes its proportional allocation of rows to the sample set of rows.

26. The apparatus of claim 23, wherein the logic for retrieving (4) further comprises logic for creating the sample set of rows using a stratified random sampling method.

45. An article of manufacture embodying logic for random sampling of rows stored in a table, wherein the table has a plurality of partitions, the method comprising: (a) determining a row count for each of the partitions of the table; (b) determining a total number of rows in the table from the row count for each of the partitions of the table; (c) computing a proportional allocation of a sample size

for each of the partitions based on the row count and the total number of rows; and (d) retrieving a sample set of rows of the sample size from the table, wherein each of the partitions of the table contributes its proportional allocation of rows to the sample set of rows.

48. The method of claim 45, wherein the retrieving (d) step further comprises creating the sample set of rows using a stratified random sampling method.

L6: Entry 3 of 5 File: USPT May 13, 2003

DOCUMENT-IDENTIFIER: US 6564221 B1
** See image for Certificate of Correction **
TITLE: Random sampling of rows in a parallel processing database system

Abstract Text (1):
A method, apparatus, and article of manufacture for random sampling of rows stored in a table, wherein the table has a plurality of partitions. A row count is determined for each of the partitions of the table and a total number of rows in the table is determined from the row count for each of the partitions of the table. A proportional allocation of a sample size is computed for each of the partitions based on the row count and the total number of rows. A sample set of rows of the sample size is retrieved from the table, wherein each of the partitions of the table contributes its proportional allocation of rows to the sample set of rows. Preferably, the computer system is a parallel processing database system, wherein each of its processing units manages a partition of the table, and some of the above steps can be performed in parallel by the processing units.

Brief Summary Text (3):
This invention relates in general to database management systems performed by computers, and in particular, to random sampling of rows in a parallel processing database system.

Brief Summary Text (9):
Thus, there is a need in the art for improved random sampling of rows stored on a database system.

Brief Summary Text (11):
The present invention discloses a method, apparatus, and article of manufacture for random sampling of rows stored in a table, wherein the table has a plurality of partitions. A row count is determined for each of the partitions of the table and a total number of rows in the table is determined from the row count for each of the partitions of the table. A proportional allocation of a sample size is computed for each of the partitions based on the row count and the total number of rows. A sample set of rows of the sample size is retrieved from the table, wherein each of the partitions of the table contributes its proportional allocation of rows to the sample set of rows. Preferably, the computer system is a parallel processing database system, wherein each of its processing units manages a partition of the table, and some of the above steps can be performed in parallel by the processing units.

Detailed Description Text (7):
FIG. 1 illustrates an exemplary hardware and software environment that could be used with the present invention. In the exemplary environment, a computer system 100 is comprised of one or more processing units (PUs) 102, also known as processors or nodes, which are interconnected by a network 104. Each of the PUs 102 is coupled to zero or more fixed and/or removable data storage units (DSUs) 106, such as disk drives, that store one or more relational databases. Further, each of the PUs 102 is coupled to zero or more data communications units (DCUs) 108, such as network interfaces, that communicate with one or more remote systems or devices.

Detailed Description Text (9):
In the preferred embodiment of the present invention, the RDBMS software comprises the Teradata.RTM. product offered by NCR Corporation, and includes one or more Parallel Database Extensions (PDEs) 112, Parsing Engines (PEs) 114, and Access Module Processors (AMPs) 116. These components of the RDBMS software perform the functions necessary to implement the RDBMS and SQL standards, i.e., definition, compilation, interpretation, optimization, database access control, database retrieval, and database update.

Detailed Description Text (12):
Both the PEs 114 and AMPs 116 are known as "virtual processors" or "vprocs". The vproc concept is accomplished by executing multiple threads or processes in a PU 102, wherein each thread or process is encapsulated within a vproc. The vproc concept adds a level of abstraction between the multi-threading of a work unit and the physical layout of the parallel processor computer system 100. Moreover, when a PU 102 itself is comprised of a plurality of processors or nodes, the vproc provides for intra-node as well as the inter-node parallelism.

Detailed Description Text (13):
The vproc concept results in better system 100 availability without undue programming overhead. The vprocs also provide a degree of location transparency, in that vprocs with each other using addresses that are vproc-specific, rather than node-specific. Further, vprocs facilitate redundancy by providing a level of isolation/abstraction between the physical node 102 and the thread or process. The result is increased system 100 utilization and fault tolerance.

Detailed Description Text (20):
The preferred embodiment of the present invention includes both a parallel sampling method and an identification of mutually exclusive rows belonging to the samples. Because the sampling method is random, any sample of the data is equally likely to be selected. In this method, however, a stratified random sampling method is used by pre-determing the number of sample rows to be obtained at each PU 102. This method still preserves most of the randomness properties as the pre-allocation does not examine the data itself. Moreover, because the rows are mutually exclusive, a row belonging to one sample can not belong to any other sample, i.e., the sampling is done without replacement.

Detailed Description Text (49):
In summary, the present invention discloses a method, apparatus, and article of manufacture for random sampling of rows stored in a table, wherein the table has a plurality of partitions. A row count is determined for each of the partitions of the table and a total number of rows in the table is determined from the row count for each of the partitions of the table. A proportional allocation of a sample size is computed for each of the partitions based on the row count and the total number of rows. A sample set of rows of the sample size is retrieved from the table, wherein each of the partitions of the table contributes its proportional allocation of rows to the sample set of rows. Preferably, the computer system is a parallel processing database system, wherein each of its processing units manages a partition of the table, and some of the above steps can be performed in parallel by the processing units.

CLAIMS:

1. A method for random sampling of rows stored in a table in a computer system, wherein the table has a plurality of partitions, the method comprising: (a) determining a row count for each of the partitions of the table; (b) determining a total number of rows in the table from the row count for each of the partitions of the table; (c) computing a proportional allocation of a sample size for each of the partitions based on the row count and the total number of rows; and (d) retrieving

a sample set of rows of the sample size from the table, wherein each of the partitions of the table contributes its proportional allocation of rows to the sample set of rows.

4. The method of claim 1, wherein the retrieving (d) step further comprises creating the sample set of rows using a stratified random sampling method.

23. An apparatus for random sampling of rows stored in a table, comprising: (a) a computer system having one or more data storage devices coupled thereto, wherein the data storage devices store at least one table, and the table has a plurality of partitions; (b) logic, performed by the computer system, for (1) determining a row count for each of the partitions of the table; (2) determining a total number of rows in the table from the row count for each of the partitions of the table; (3) computing a proportional allocation of a sample size for each of the partitions based on the row count and the total number of rows; and (4) retrieving a sample set of rows of the sample size from the table, wherein each of the partitions of the table contributes its proportional allocation of rows to the sample set of rows.

26. The apparatus of claim 23, wherein the logic for retrieving (4) further comprises logic for creating the sample set of rows using a stratified random sampling method.

45. An article of manufacture embodying logic for random sampling of rows stored in a table, wherein the table has a plurality of partitions, the method comprising: (a) determining a row count for each of the partitions of the table; (b) determining a total number of rows in the table from the row count for each of the partitions of the table; (c) computing a proportional allocation of a sample size for each of the partitions based on the row count and the total number of rows; and (d) retrieving a sample set of rows of the sample size from the table, wherein each of the partitions of the table contributes its proportional allocation of rows to the sample set of rows.

48. The method of claim 45, wherein the retrieving (d) step further comprises creating the sample set of rows using a stratified random sampling method.